arXiv:1410.7241v1 [stat.ML] 27 Oct 2014

A Greedy Homotopy Method for Regression with Nonconvex Constraints

Fabian L. Wauthier
flw@well.ox.ac.uk

Peter Donnelly
donnelly@well.ox.ac.uk

WTCHG and Department of Statistics
University of Oxford

24th of October 2014

Abstract

Constrained least squares regression is an essential tool for high-dimensional data analysis. Given a partition $\mathcal{G}$ of input variables, this paper considers a particular class of nonconvex constraint functions that encourage the linear model to select a small number of variables from a small number of groups in $\mathcal{G}$. Such constraints are relevant in many practical applications, such as Genome-Wide Association Studies (GWAS). Motivated by the efficiency of the Lasso homotopy method [3, 14], we present RepLasso, a greedy homotopy algorithm that tries to solve the induced sequence of nonconvex problems by solving a sequence of suitably *adapted* convex surrogate problems. We prove that in some situations RepLasso recovers the global minima of the nonconvex problem. Moreover, even if it does not recover global minima, we prove that in relevant cases it will still do no worse than the Lasso in terms of support and signed support recovery, while in practice outperforming it. We show empirically that the strategy can also be used to improve over other Lasso-style algorithms. Finally, a GWAS of ankylosing spondylitis highlights our method's practical utility.

1 Introduction

We are interested in model parsimony in the context of linear observation models of the form

$$y = X\beta^* + w \qquad w \sim \mathcal{N}(0, \sigma^2 I), \tag{1}$$

where X is an $n \times p$ matrix of covariates, β^* is a regression parameter, w is a noise vector and y is a vector of responses. Given X, y, a constraint function $\Omega(\cdot)$, and constraint parameter $\tau > 0$, constrained least squares regression estimates β^* as

$$\operatorname*{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y - X\beta\|_2^2 \quad \text{s.t.} \quad \Omega(\beta) \leq \tau. \tag{G}$$

A closely related formulation writes (G) in penalized form using penalizer $\Omega(\cdot)$ and penalty parameter $\lambda > 0$. In the following, we will motivate our algorithm using the constrained formulation (G). However, the proposed algorithm is then more naturally expressed as solving a sequence of penalized problems.

The use of sparsity for model selection is an integral component of the modern statistics toolbox and is especially relevant for the $n \ll p$ case. The Lasso [17] is a well-known special case of (G) which replaces a hard ℓ_0 constraint (i.e., $\|\beta\|_0 \leq \tau$) by an ℓ_1 surrogate (i.e., $\|\beta\|_1 \leq \tau$) that retains some sparsity-inducing properties. Since the Lasso regularization path is continuous and

piecewise linear [16], it can be easily traced out using the homotopy method [3, 14]. This is done by writing (G) in penalized form and then tracing out $\lambda = \infty \downarrow 0$. In the following, we will often refer to the Lasso homotopy method simply as the Lasso. The efficiency of the homotopy method is one of the main benefits of the ℓ_1 relaxation approach and is key for efficient model selection.

Recently, there has been increased interest in enhancing the Lasso with structured sparsity. Relevant examples include [8, 10, 11, 22, 25]. In all these cases the structured sparsity is induced by replacing the ℓ_1 penalizer of the Lasso with more complex, yet still convex, penalizers. Because the overall objective remains convex in β, efficient algorithms exist to solve these problems. While these methods have many practical applications, the focus on convex formulations has necessarily excluded important inference problems that cannot easily be phrased in terms of structured convex objectives. This paper was motivated by applications where for a given partition $\mathcal{G}$ of variables $\{1, \ldots, p\}$ it is reasonable to assume that in *at most* a few groups *at most* a few *representatives* are relevant for predicting y. For instance, in the Genome-Wide Association Study (GWAS) we consider in Section 6, it is reasonable to suppose that in at most a small number of genes at most a small number of SNPs are associated with the response variable. To suite these and other applications, we are interested in constraint functions $\Omega(\cdot)$ that encourage the estimate of β^* to be nonzero on at most a few elements in at most a few groups $G \in \mathcal{G}$, which can be thought of as orthogonal to the Group Lasso [22]. This problem is not adequately solved by the Exclusive/Elitist Lasso [11, 25], which is a convex formulation that generally selects *at least one* but at most a few variables from each group. This is problematic if we believe that most groups in $\mathcal{G}$ will contain no relevant variables for predicting y, as for example in the GWAS application. To the best of our knowledge the selection behavior we seek can only be achieved by nonconvex constraints.

Given a parameter $\theta \in \mathbb{R}^p$ and the partition $\mathcal{G}$, we will encode our constraints as a nonconvex function $\Omega_{\theta,\mathcal{G}}(\cdot)$. For fixed penalty parameter $\lambda > 0$, there are specialized methods for the nonconvex penalized cousin of (G) (e.g., [2, 4, 7, 13, 28]). While these methods might be appropriate for finding a local minimum of (G) with $\Omega_{\theta,\mathcal{G}}(\cdot)$ and some τ fixed, they are not useful for developing a homotopy-like algorithm which allows τ to range over an interval. Motivated by the practicality and efficiency of the Lasso homotopy method, we propose RepLasso (for "Representative Lasso"), a homotopy-like algorithm that attempts to fill this gap. At a high level, RepLasso tries to build and solve a sequence of convex surrogates so that, as τ is swept out, the boundary of the surrogate constraint ball locally approximates the boundary of the ball induced by $\Omega_{\theta,\mathcal{G}}(\cdot)$. A crucial feature that allows us to do this efficiently is that the nonconvex constraint balls induced by $\Omega_{\theta,\mathcal{G}}(\cdot)$ can be decomposed as unions of convex balls. Moreover, the sequence of surrogates is chosen so that the induced regularization path is continuous and piecewise linear and can thus be efficiently traced out using a homotopy-style algorithm.

To motivate the algorithm, we show theoretically that, under certain conditions, RepLasso traces out the global minima of (G) with constraint function $\Omega_{\theta,\mathcal{G}}(\cdot)$. More importantly, we prove that, even though RepLasso may not exactly solve this problem in general, on relevant problems it will still do at least as well as the Lasso in terms of support subset and signed support recovery. In practice, a strict improvement is observed. A class of Lasso-style algorithms has recently been popularized which pre-process X, y in some way, prior to solving a standard Lasso problem (e.g., [6, 9, 15, 26]). As we demonstrate in Section 6, RepLasso can also yield strict improvements in these settings. Furthermore, RepLasso can be usefully applied to ℓ_1 constrained logistic regression [12], as we demonstrate in a GWAS application. Lastly, we prove in the Supplementary Material that, given some mild assumptions, a variant of RepLasso cannot do worse than the well-known Lars algorithm of Efron et al. [3].

The paper is organized as follows: We review related research in Section 2 before introducing $\Omega_{\theta,\mathcal{G}}(\cdot)$ and simplifying (G) in Section 3. In Section 4 we present the RepLasso as a generalization of the Lasso homotopy method and in Section 5 give a theoretical comparison of Lasso and RepLasso. Results on synthetic data and a GWAS application are given in Section 6. We conclude with final remarks in Section 7. Proofs are collected in the Supplementary Material.

2 Related Research

Nonconvex penalties for least squares regression have been (for example) considered by Fan and Li [4] and Zhang [23]. Methods for optimizing convex loss functions with nonconvex regularizers include, among others, local quadratic approximation [4], minorization-maximization [7], local linear approximation [28] and composite gradient descent [13]. The Adaptive Lasso of [2] is also related to our method. However, a drawback of all of these approaches is that they focus on a single optimization problem, indexed by a fixed penalty parameter. This precludes their use for efficiently minimizing a sequence of problems (G) indexed by τ, as in a homotopy method. The various applications of the homotopy idea have so far focused on other convex problems. Well-known examples are the Elastic Net [27] and the SVM [5]. However, there are very few extensions to nonconvex least squares problems. One of the few methods that efficiently sweeps out local minima paths of such problems is due to Zhang [23]. However, as [23] assumes the penalty to be separable across the p coefficients, it is not useful for the type of structured sparsity we consider in this paper. There has been growing interest in more complex sparsity patterns induced by structured penalties. Among convex extensions, the Group Lasso [22] is a well-known example, which for some partition $\mathcal{G}$, replaces the ℓ_1 penalty above by a sum of ℓ_2 penalties over groups of variables indexed by $G \in \mathcal{G}$. This method will select groups of variables, not representatives, and so can be seen as a counterpart to the work in this paper. Several variations of this approach have been proposed [8, 10]. The Exclusive/Elitist Lasso [11, 25] is more closely aligned with our goal. However, this method effectively encourages each group to contribute at least one variable to the support set. In contrast, our method encourages the selection of a small number of variables in a small number of groups. Finally, while there are other structured, nonconvex penalizers (e.g., [20]), there are no homotopy algorithms to solve them.

3 Structured Nonconvex Problems

Many situations exist where for some partition $\mathcal{G}$ of $\{1, \ldots, p\}$ we know that β^* contains at most a few nonzero elements in at most a few groups $G \in \mathcal{G}$.[1] Given a partition $\mathcal{G} = \{G_1, \ldots, G_g\}$ without singleton or empty sets, and a vector $\theta = (\theta_1, \ldots, \theta_g) \geq 0$, the following constraint function targets this situation

$$\Omega_{\theta,\mathcal{G}}(\beta) = \sum_{i<j \in G_{g'} \in \mathcal{G}} \frac{\omega_{\theta_{g'}}(\beta_i, \beta_j)}{|G_{g'}| - 1} \tag{2}$$

$$\omega_{\theta_{g'}}(\beta_i, \beta_j) = \min(|\beta_i|, |\beta_j|)(1 + \theta_{g'}) + \max(|\beta_i|, |\beta_j|).$$

Let $B_{\theta,\mathcal{G}}(\tau) = \{\beta \in \mathbb{R}^p : \Omega_{\theta,\mathcal{G}}(\beta) \leq \tau\}$ be the induced constraint balls. We are interested in the following nonconvex instance of (G) with the constrained objective $J_\tau(\beta)$ over β, indexed by τ

$$\beta(\tau) \in \underset{\beta \in \mathbb{R}^p}{\operatorname{argmin}} J_\tau(\beta) = \underset{\beta \in \mathbb{R}^p}{\operatorname{argmin}} \begin{cases} \frac{1}{2n} \|y - X\beta\|_2^2 & \text{if } \beta \in B_{\theta,\mathcal{G}}(\tau) \\ \infty & \text{o.w.} \end{cases}. \tag{P1}$$

If $\theta = \mathbf{0}$, then $\Omega_{\theta,\mathcal{G}}(\beta) = \|\beta\|_1$ for all $\mathcal{G}$, and so $(P1)$ recovers the Lasso problem as special case. However, when $\theta \neq \mathbf{0}$, the constraint function is non-separable, non-convex and induces star-shaped balls, as exemplified in Figure 1. Note that when a subset of components of θ is set to zero, we can effectively treat the variables corresponding to those groups as ungrouped, as they only contribute an ℓ_1 penalty to $\Omega_{\theta,\mathcal{G}}(\beta)$. When $\theta \neq \mathbf{0}$, $\Omega_{\theta,\mathcal{G}}(\beta)$ can be thought of as a nonconvex counterpart to the well-known Group Lasso penalty [22], where the nonconvexity encourages solutions $\beta(\tau)$ of $(P1)$ to select *at most* a small number of *representatives* from *at most* a few groups $G \in \mathcal{G}$. The penalty $\Omega_{\theta,\mathcal{G}}(\beta)$ is also distinct from the convex Exclusive/Elitist Lasso penalty [11, 25], which effectively encourages each group to select at least one variable.

[1] In the GWAS application in Section 6, $\mathcal{G}$ corresponds to a partition of SNPs by genes and we know that in at most a few genes at most a few SNPs are truly relevant for predicting y.

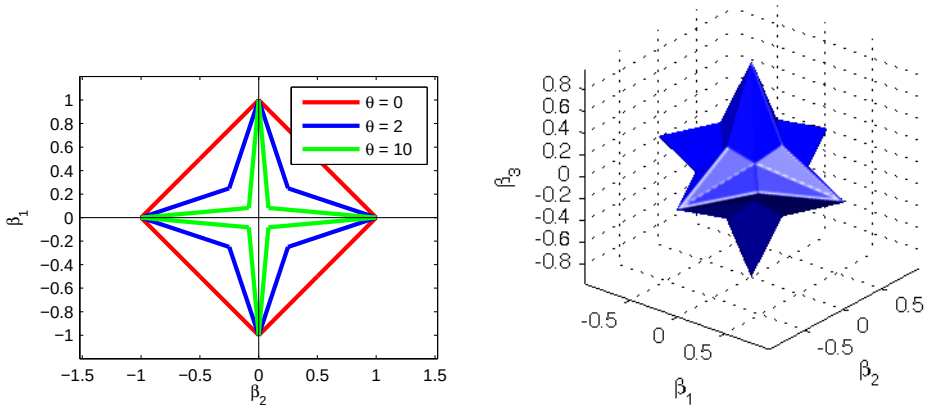

(a) (b)

Figure 1: (a) The balls induced by Eq. (2) for $\mathcal{G} = \{\{1, 2\}\}$ and varying values of θ. If $\theta = 0$ we recover the ℓ_1 norm as a special case. (b) A ball for $\mathcal{G} = \{\{1, 2, 3\}\}, \theta = 2$.

For some positive vector s, let $B_s(\tau) = \{\beta \in \mathbb{R}^p : \|\text{diag}(s)\beta\|_1 \leq \tau\}$ be the s-weighted ℓ_1 ball. An important property that is suggested by Figure 1 is that $B_{\theta,\mathcal{G}}(\tau)$ can be written as a union of weighted ℓ_1 balls and so has planar faces. Let $\Gamma(i) \in \{1, \ldots, g\}$ be the (unique) variable index so that $i \in G_{\Gamma(i)}$.

Proposition 1 (Union Decomposition)**.** *Let the partition be $\mathcal{G} = \{G_1, \ldots, G_g\}$ and the parameter $\theta = (\theta_1, \ldots, \theta_g) \geq 0$. There is a finite set $\mathcal{S}_{\theta,\mathcal{G}} \subset \mathbb{R}^p$ of vectors $s \geq \mathbf{1}$, so that for any $\tau > 0$*

$$B_{\theta,\mathcal{G}}(\tau) = \bigcup_{s \in \mathcal{S}_{\theta,\mathcal{G}}} B_s(\tau). \tag{3}$$

Define $\Pi_{g'}$ to be all permutations $\pi_{g'}$ of the elements in $G_{g'}$ and let $\Pi_{\mathcal{G}} = \times_{g'=1}^{g} \Pi_{g'}$ be their cross-product, whose elements $\pi \in \Pi_{\mathcal{G}}$ are g-tuples of permutations $\pi = (\pi_1, \ldots, \pi_g)$. For some $\pi \in \Pi_{\mathcal{G}}$, denote by $\pi_{\Gamma(i)}(i) \in \{1, \ldots, |G_{\Gamma(i)}|\}$ the position of $i \in G_{\Gamma(i)}$ in permutation $\pi_{\Gamma(i)}$. We have

$$\mathcal{S}_{\theta,\mathcal{G}} = \cup_{\pi \in \Pi_{\mathcal{G}}} \{s_\pi\} \tag{4}$$

$$s_{\pi,i} = 1 + (\pi_{\Gamma(i)}(i) - 1)\frac{\theta_{\Gamma(i)}}{|G_{\Gamma(i)}| - 1}. \tag{5}$$

A common, brute force approach that eliminates the computational issues of $(P1)$ replaces $B_{\theta,\mathcal{G}}(\tau)$ by its convex hull. However, in this case the convex hull is the ℓ_1 ball $B_1(\tau)$ which eliminates all structural information inherited from $\theta, \mathcal{G}$ and would not lead to the desired selection behavior. In this paper we advocate an orthogonal strategy that instead focuses on replacing $B_{\theta,\mathcal{G}}(\tau)$ by a suitable *sequence* of *weighted* ℓ_1 balls, indexed by τ. To achieve this, we will exploit the decompositional structure of $B_{\theta,\mathcal{G}}(\tau)$ highlighted in Proposition 1. Specifically, our method is motivated by the following extension of a well-known result of Rosset and Zhu [16].

Proposition 2 (Local Piecewise Linearity)**.** *Suppose X has absolutely continuous distribution and that $\exists \tau' > 0$ s.t. $\nexists \beta \in B_{\theta,\mathcal{G}}(\tau')$ which is a minimum of $\|y - X\beta\|_2^2$. Let $\tau_{\max}$ be the supremum over these τ'. The set of local minima of $J_\tau(\cdot)$ in $(P1)$ with $\tau \in (0, \tau_{\max})$ is with probability 1 a finite union of piecewise linear paths, each path indexed by τ and lying on a ball $B_s(\tau), s \in \mathcal{S}_{\theta,\mathcal{G}}$.*

Proposition 2 emphasizes that for the range of interesting values of $\tau \in (0, \tau_{\max})$, the local minima of $J_\tau(\cdot)$ in $(P1)$ can be grouped into a set of local minima *paths*, each indexed by τ. Moreover,

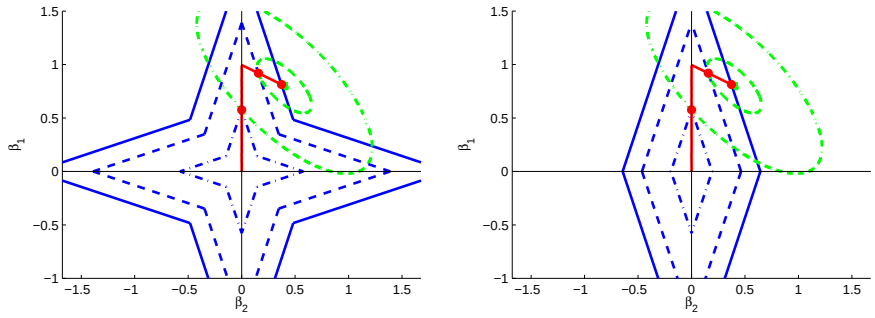

(a) (b)

Figure 2: (a) The global minima path for $B_{2,\{\{1,2\}\}}(\tau)$. The dashed-dotted, dashed and solid constraint boundaries correspond to progressively larger τ. The solutions on the regularization path (red) are defined by intersections of the constraint set with the corresponding ellipsoid (red dots). (b) The corresponding regularization path for the weighted ℓ_1 ball $B_{(1,3)^\top}(\tau)$. If the weighting is known the same regularization path can be reproduced.

any such local minimum path lies on some weighted ℓ_1 ball $B_s(\tau)$, with $s \in \mathcal{S}_{\theta,\mathcal{G}}$ appropriately chosen. With the aid of Proposition 2, it is possible to re-express $(P1)$ as a special set of penalized optimization problems, indexed by τ. This change of representation will be useful for the homotopy-like algorithm we present shortly. By Proposition 2 and convexity, for any solution $\beta(\tau)$ of $(P1)$ with $\tau \in (0, \tau_{\max})$ (i.e., a global minimum of $J_\tau(\cdot)$), $\exists \lambda^*(\tau), s^*(\tau) \in \mathcal{S}_{\theta,\mathcal{G}}$ so that

$$\beta(\tau) \in \underset{\beta \in \mathbb{R}^p}{\operatorname{argmin}} \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda^*(\tau) \|\operatorname{diag}(s^*(\tau))\beta\|_1. \tag{P2}$$

Thus, modulo uniqueness issues, there exist $\lambda^*(\tau)$, $s^*(\tau)$ so that $(P2)$ is in some sense equivalent to $(P1)$. Figure 2 shows a motivating example of this. In this case, the global minimum path of Figure 2(a) could be reproduced using the $B_{(1,3)^\top}(\tau)$ balls of Figure 2(b). Of course, knowledge of the vector-valued function $s^*(\tau) \in \mathcal{S}_{\theta,\mathcal{G}}$ would imply knowing for each τ roughly where on $B_{\theta,\mathcal{G}}(\tau)$ the global minimum of $J_\tau(\cdot)$ in $(P1)$ lies, which is hard in general. We thus cannot expect to be able to efficiently produce the *entire* regularization path of $(P1)$ for all $\tau \in (0, \tau_{\max})$ using the equivalence between $(P2)$ and $(P1)$.

A Simplifying Assumption. The formulation in $(P2)$ replicates the regularizing effect of $B_{\theta,\mathcal{G}}(\tau)$ in $(P1)$ using a sequence of weighted ℓ_1 balls that depend on τ (characterized by $s^*(\tau)$). This dependence is necessary as the global minimum of $J_\tau(\cdot)$ in $(P1)$ can "jump" from one weighted ℓ_1 ball to another as we vary $\tau \in (0, \tau_{\max})$. If we let $S(\beta)$ be the support of some vector β, we can simplify the problem of finding sequences $\lambda^*(\tau), s^*(\tau)$ for $(P2)$, by assuming that

A0: $\exists s^* \in \mathcal{S}_{\theta,\mathcal{G}}$ so that $\forall \tau \in (0, \tau_{\max})$, $(P1)$ has a unique solution which lies in $B_{s^*}(\tau)$. For any $0 < \tau_1 < \tau_2 < \tau_{\max}$, the solutions of $(P1)$ satisfy $S(\beta(\tau_1)) \subseteq S(\beta(\tau_2))$.

Under **A0**, the problem immediately reduces to finding a sequence $\lambda^*(\tau)$ and a single positive vector $s^* \in \mathcal{S}_{\theta,\mathcal{G}}$. In fact, it is not even necessary to know the precise function $\lambda^*(\tau)$: For any $\lambda > 0$, so long as the solution $\bar{\beta}(\lambda)$ to $(P2)$ with $\lambda^*(\tau)$ replaced by λ and $s^*(\tau) = s^*$, satisfies for $\tau \triangleq \left\|\operatorname{diag}(s^*)\bar{\beta}(\lambda)\right\|_1$ that $\tau \in (0, \tau_{\max})$, we know that $\lambda = \lambda^*(\tau)$. Thus, under **A0** we only seek to find the vector $s^* \in \mathcal{S}_{\theta,\mathcal{G}}$ so that solving $(P1)$ is for some λ equivalent to solving

$$\bar{\beta}(\lambda) \in \underset{\beta \in \mathbb{R}^p}{\operatorname{argmin}} \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda \|\operatorname{diag}(s^*)\beta\|_1. \tag{S}$$

Algorithm 1: REPLASSO$(X, y, \mathcal{G}, \theta)$

$\bar{y} = 0, A = (), L = 0, \lambda = \|X^\top y\|_\infty, s(\lambda) = \mathbf{1}, \bar{\beta}(\lambda) = 0$
while $\lambda > 0$

Stage 1
- **if** $L = 0$ # Add a variable
  - $A = (A, i^*)$, where $i^* = \text{argmax}_{j \in A^c} \left| X_j^\top (y - \bar{y}) / s_j(\lambda) \right|$
  - $s_M(\lambda^-) = s_M(\lambda) + \frac{\theta_{\Gamma(i^*)}}{|G_{\Gamma(i^*)}| - 1}\mathbf{1}, \quad s_{M^c}(\lambda^-) = s_{M^c}(\lambda),$ with $M = \left\{A^c \cap G_{\Gamma(i^*)}\right\}$
- **if** $L = 1$ # Delete a variable
  - $A = A \backslash i^*$, where $i^* = \arg_{i \in A} [\![\bar{\beta}_i(\lambda) = 0]\!]$

Stage 2
- $\bar{w}_A = A_A \left(X_A^\top X_A\right)^{-1} \text{diag}(\text{sgn}\left(X_A^\top (y - \bar{y})\right)) s_A(\lambda)$, with A_A s.t. $\|X_A \bar{w}_A\|_2^2 = 1$

Stage 3
- Find smallest $\rho > 0$ s.t.
  - $\exists j \in A^c$ s.t. $|X_j^\top (y - \bar{y} - \rho X_A \bar{w}_A)/s_j(\lambda)| = \lambda - \rho$: set $L = 0$
  - $\exists i \in A$ s.t. $\bar{\beta}_i(\lambda) \neq 0$ and $\bar{\beta}_i(\lambda) + \rho w_i = 0$: set $L = 1$

Stage 4
- $\bar{\beta}_A(\lambda - \rho) = \bar{\beta}_A(\lambda) + \rho \bar{w}_A, \quad \bar{\beta}_{A^c}(\lambda - \rho) = 0, \quad \bar{y} = X\bar{\beta}(\lambda - \rho)$
- $\lambda = \lambda - \rho$

return $\bar{\beta}$

This paper makes two main contributions. The first contribution in Section 4 proves that if **A0** holds, then there is an algorithm, RepLasso, which (effectively) greedily estimates the vector s^* making (S) and $(P1)$ equivalent while sweeping out $\lambda > 0$ and producing solutions $\bar{\beta}(\lambda)$ in a homotopy-like fashion. Of course, if **A0** does not hold, there may not be an equivalence between (S) and $(P1)$. In that case, we may think of (S) as a convex surrogate for $(P1)$ for some vector s^* that is greedily constructed by RepLasso. The second contribution of this paper is to prove in Section 5 that, whether **A0** holds or not, RepLasso will in relevant regression problems still perform at least as well as the Lasso in terms of variable selection. Empirical evidence in Section 6 shows a strict improvement in practice.

4 RepLasso: A Greedy Homotopy Method

To motivate our description of RepLasso, we first make the following observation regarding the sensitivity of problems in the form of (S) to approximations of s^*. For a positive vector b, let $\bar{\beta}_b(\lambda)$ be a solution to (S) with penalty $\lambda \|\text{diag}(b)\beta\|_1$.

Proposition 3 (Recoverability of (S))**.** *Suppose X has absolutely continuous distribution. For any vectors $a \geq b \geq \mathbf{1}$ and $\lambda > 0$, with probability 1 $\bar{\beta}_a(\lambda), \bar{\beta}_b(\lambda)$ are unique. If additionally $\|diag(a)\bar{\beta}_b(\lambda)\|_1 = \|diag(b)\bar{\beta}_b(\lambda)\|_1$, then $\bar{\beta}_a(\lambda) = \bar{\beta}_b(\lambda)$.*

Thus, if $\bar{\beta}_b(\lambda)$ has zero coefficients, then it doesn't matter if on those coefficients b underestimates the value of a, so long as b matches a on the remaining coefficients.

The RepLasso algorithm (Algorithm 1) is a generalization of the Lasso homotopy method [3, 14] that exploits Proposition 3 to solve (S). If X is absolutely continuous and **A0** holds, then Proposition 3 suggests the existence of a sequence $s(\lambda)$, satisfying $\forall \lambda > 0, s^* \geq s(\lambda) \geq \mathbf{1}$, so that with probability 1 (S) can $\forall \lambda > 0$ be solved as $\bar{\beta}(\lambda) \triangleq \bar{\beta}_{s^*}(\lambda) = \bar{\beta}_{s(\lambda)}(\lambda)$. As Theorem 1 shows, RepLasso computes such a sequence $s(\lambda)$, while simultaneously producing solutions $\bar{\beta}_{s(\lambda)}(\lambda)$. Notice that RepLasso is identical to the Lasso homotopy method if $\theta = \mathbf{0}$ (which means that $\forall \lambda > 0, s(\lambda) = \mathbf{1}$). The only differences are that $s(\lambda) \neq \mathbf{1}$ when $\theta \neq \mathbf{0}$. We will discuss RepLasso as constructive proof for Theorem 1.

Theorem 1 (RepLasso)**.** *Assume that X has absolutely continuous distribution and that* **A0** *holds. Let $s^* \in \mathcal{S}_{\theta,\mathcal{G}}$ be the vector so that $(P1)$ is equivalent to (S). Then with probability* 1*, RepLasso produces a sequence $s(\lambda)$ so that $\bar{\beta}_{s^*}(\lambda) = \bar{\beta}_{s(\lambda)}(\lambda)$. By the equivalence of $(P1)$ and (S), it follows that with probability 1, RepLasso produces the global minima of $(P1)$.*

Proof. Note from our earlier discussion that it is sufficient for RepLasso to estimate sequences $s(\lambda)$ which are *piecewise constant* with changepoints at values λ_t where the support of $\bar{\beta}_{s^*}(\lambda_t)$ changes. By **A0**, we know that the support of $\bar{\beta}_{s^*}(\lambda)$ is monotonically increasing with λ decreasing. Hence, we only need to discuss the variable addition case (case $L = 0$ in stage 1) of RepLasso for this argument. Conceptually, RepLasso first initializes $s(\infty) = \mathbf{1}$ (for practical reasons it suffices to start at $\lambda = \|y^\top X\|_\infty$). Then, while keeping $s(\lambda) = s(\infty)$ constant, RepLasso (conceptually) traces out $\lambda = \infty \downarrow 0$ while solving $\bar{\beta}_{s(\lambda)}(\lambda) = \mathbf{0}$ until reaching $\lambda_1 = \|y^\top X\|_\infty$, where the first variable i_1^* is selected by $\bar{\beta}_{s(\lambda)}(\lambda)$ (the $L = 0$ case in stage 1 of RepLasso). Because $s(\lambda) = \mathbf{1}$ was up to now fixed, RepLasso is up to this point identical to the Lasso homotopy method. Due to Proposition 3, we know that with probability 1, $\forall \lambda \in [\lambda_1, \infty]$ we have $\bar{\beta}_{s^*}(\lambda) = \bar{\beta}_{s(\lambda)}(\lambda)$. Under **A0**, we know that $\forall 0 < \lambda \leq \lambda_1$, i_1^* will remain selected and that the relative order of i_1^* in the set of variables $G_{\Gamma(i_1^*)}$, as induced by the magnitude of their coefficients in $\bar{\beta}_{s(\lambda_1)}(\lambda_1)$ will not change. Using this and the general form of $s^* \in \mathcal{S}_{\theta,\mathcal{G}}$ given by Proposition 1, we can modify $s(\lambda)$ in a way that is consistent with Proposition 3. Specifically, if we let $t = 1$, then the current active set is $A_t = \{i : \left|X_i^\top (y - X\bar{\beta}_{s(\lambda_t)}(\lambda_t))\right| / s_i(\lambda_t) = \lambda_t\}$. We may apply the following generic update to $s(\lambda)$ so that at λ_t^- (i.e., for a value of λ infinitesimally smaller than λ_t) it satisfies

$$s_j(\lambda_t^-) = \begin{cases} s_j(\lambda_t) + \frac{\theta_{\Gamma(i_t^*)}}{|G_{\Gamma(i_t^*)}| - 1} & j \in \{A_t^c \cap G_{\Gamma(i_t^*)}\} \\ s_j(\lambda_t) & \text{o.w.} \end{cases}. \tag{6}$$

Notice that the change leaves the path $\bar{\beta}_{s(\lambda_t)}(\lambda_t)$ continuous in the neighborhood of λ_t. RepLasso then continues to decrease $\lambda = \lambda_1 \downarrow 0$, again keeping $s(\lambda) = s(\lambda_1^-)$ constant and producing solutions $\bar{\beta}_{s(\lambda)}(\lambda)$ along the way, until a point $\lambda_2 > 0$ is reached when a new variable is selected by $\bar{\beta}_{s(\lambda)}(\lambda)$. Because $s(\lambda)$ was kept constant for $\lambda \in [\lambda_2, \lambda_1^-]$, this can be achieved by a straightforward modification of the Lasso homotopy method[2]. As before, we know from our update of $s(\lambda)$ and Proposition 3 that with probability 1, $\forall \lambda \in [\lambda_2, \lambda_1]$ we have $\bar{\beta}_{s^*}(\lambda) = \bar{\beta}_{s(\lambda)}(\lambda)$. At this point, **A0** and Proposition 1 again allow us to update $s(\lambda)$ using Eq. (6) with $t = 2$. RepLasso continues sweeping out λ in this fashion until some final value $\lambda_T > 0$ is reached. By the time the algorithm has completed, we know that with probability 1, $\forall \lambda \in [\lambda_T, \infty]$ we have $\bar{\beta}_{s^*}(\lambda) = \bar{\beta}_{s(\lambda)}(\lambda)$. The final claim follows immediately. □

When **A0** does not hold, we can apply Proposition 3 to $(P2)$ to see that RepLasso will generally still recover global minima of $(P1)$ for large $\lambda > 0$. Indeed, if RepLasso adds variables one by one, the first variable selected by RepLasso is also the first selected by $(P1)$. Regardless of whether **A0** holds, Section 5 shows strong results for RepLasso relative to the ℓ_1 relaxation of $(P1)$ (i.e., the Lasso).

5 Comparing RepLasso and Lasso

In this section we show several results irrespective of whether **A0** holds, but assuming that $\mathcal{G}$, β^* satisfy certain conditions. Before continuing, we briefly outline some notation. Let X_j be the column j of X and X_A a matrix which consists of the columns indexed by A. Let the support set of β^* be $S \triangleq S(\beta^*)$. Denote the signed support of β^* by $S_\pm = S_\pm(\beta^*)$, where

$$S_\pm(\beta_i) \triangleq \begin{cases} +1 & \text{if } \beta_i > 0 \\ -1 & \text{if } \beta_i < 0 \\ 0 & \text{o.w.} \end{cases}. \tag{7}$$

Our analysis in this section relies on various subsets of the following assumptions

[2]Specifically, where the Lasso homotopy method traces out *equiangular* directions, the RepLasso follows *skew-angular* directions (given in stage 2), with the angle skew determined by the weights $s_A(\lambda)$.

A1: $\forall G \in \mathcal{G}, |\{i \in G : \beta_i^* \neq 0\}| \leq 1$

A2: $\forall A \subset S$ and u_A the equiangular vector in Eq. (2.6) of [3], $\nexists j \in A^c, |X_A^\top u_A| = |X_j^\top u_A|\mathbf{1}$

A3: $X_S^\top X_S$ is invertible

A4: Following [19, 24], define

$$\mu_j \triangleq X_j^\top X_S (X_S^\top X_S)^{-1} \mathrm{sgn}(\beta_S^*)$$
$$\gamma_i \triangleq e_i^\top \left(X_S^\top X_S\right)^{-1} n.$$

We have $\forall j \in S^c, |\mu_j| < 1$, $\forall i \in S, \mathrm{sgn}(\beta_i^*)\gamma_i > 0$.

Assumption **A1** formalizes that β^* is nonzero on at most a few elements (in this case one) of each group of $\mathcal{G}$. Assumption **A2** ensures that the active set estimated by the algorithm in [3] matches the support set and holds with probability 1 if X has spherical and absolutely continuous distribution. Assumption **A3** is a zeroth-order condition necessary for identifiability. Assumption **A4** ensures that the Lasso has any chance of recovering the signed support of β^*.

Many analyses of the Lasso focus on its support recovery properties. The following theorem compares RepLasso against Lasso in terms of this measure.

Theorem 2 (Support Subset Recovery)**.** *Assume that* **A1**–**2** *hold. Denote by* $\hat{\beta}(\lambda)$ *and* $\bar{\beta}(\lambda)$ *the Lasso and RepLasso solutions for penalty parameter* λ*. Given* X, y*, we have for any* $\lambda_{\min} > 0$

$$\forall \lambda \geq \lambda_{\min}\ S(\hat{\beta}(\lambda)) \subseteq S \implies \forall \lambda \geq \lambda_{\min}\ S(\bar{\beta}(\lambda)) \subseteq S.$$

Proof sketch. Suppose $\forall \lambda \geq \lambda_{\min}\ S(\hat{\beta}(\lambda)) \subseteq S$. For $t' \leq t$, let $\hat{A}_{t'}$ and $\bar{A}_{t'}$ be the active sets of the Lasso and RepLasso at iteration t' until $\lambda_{\min}$ is reached. By **A2** we can show that $\forall t' \leq t, \hat{A}_{t'} \subseteq S$. Since $s(\|X^\top y\|_\infty) = \mathbf{1}$ we know that $\hat{A}_1 = \bar{A}_1 \subseteq S$. It then follows from $\hat{A}_2 \subseteq S$, assumption **A1** and the construction of $s(\lambda)$ that we also have $\hat{A}_2 = \bar{A}_2 \subseteq S$. Iterating this argument over $t' \leq t$, we can then show that $\forall t' \leq t, \bar{A}_{t'} \subseteq S$ from which it follows that $\forall \lambda \geq \lambda_{\min}\ S(\bar{\beta}(\lambda)) \subseteq S$.

In many situations we are not only interested in recovering a subset of the true support, but the signed support of β^*. The previous result can be strengthened to cover this case.

Theorem 3 (Signed Support Recovery)**.** *Assume that* **A1**–**4** *hold. Denote by* $\hat{\beta}(\lambda)$ *and* $\bar{\beta}(\lambda)$ *the Lasso and RepLasso solutions using penalty parameter* λ*. For any* $\lambda_{\min} > 0$*, we have with probability 1 over an absolutely continuous distribution on noise* w

$$\forall \lambda \geq \lambda_{\min}\ S(\hat{\beta}(\lambda)) \subseteq S, S_\pm(\hat{\beta}(\lambda_{\min})) = S_\pm \implies \forall \lambda \geq \lambda_{\min}\ S(\bar{\beta}(\lambda)) \subseteq S, S_\pm(\bar{\beta}(\lambda_{\min})) = S_\pm.$$

Proof sketch. Suppose $\forall \lambda \geq \lambda_{\min}\ S(\hat{\beta}(\lambda)) \subseteq S, S_\pm(\hat{\beta}(\lambda_{\min})) = S_\pm$. For $t' \leq t$, let $\hat{A}_{t'}, \bar{A}_{t'}$ be the active sets of Lasso/RepLasso until $\lambda_{\min}$ is reached. As in Theorem 2, we conclude from **A2** that $\forall t' \leq t$, $\hat{A}_{t'} \subseteq S$, which then tells us via **A1** that $\forall t' \leq t, \bar{A}_{t'} \subseteq S$ and so $\forall \lambda \geq \lambda_{\min}, S(\bar{\beta}(\lambda)) \subseteq S$. Furthermore, by **A1** and the construction of $s(\lambda)$, we know $\forall \lambda \geq \lambda_{\min}$ that $s_S(\lambda) = \mathbf{1}$ and $s_{S^c}(\lambda) \geq \mathbf{1}$. Utilizing **A3**–**4** and Lemma 1 of Wauthier et al. [21] (which holds with probability 1 over noise w), we can then argue that with probability 1, $S_\pm(\bar{\beta}(\lambda_{\min})) = S_\pm$.

Consequences for other Methods. Besides the Lasso, Theorems 2 and 3 also apply to many related algorithms that pre-process the data X, y in some way, prior to running the Lasso on the modified data. Instances of these algorithms are, for example, the Adaptive Lasso [26] and various Preconditioned Lasso algorithms [6, 9, 15]. Indeed, if the relevant assumptions **A1**–**4** hold, the result is even true for ℓ_1 regularized minimization of quadratic approximations to logistic regression as proposed in [12]. We will empirically highlight this property in Section 6.

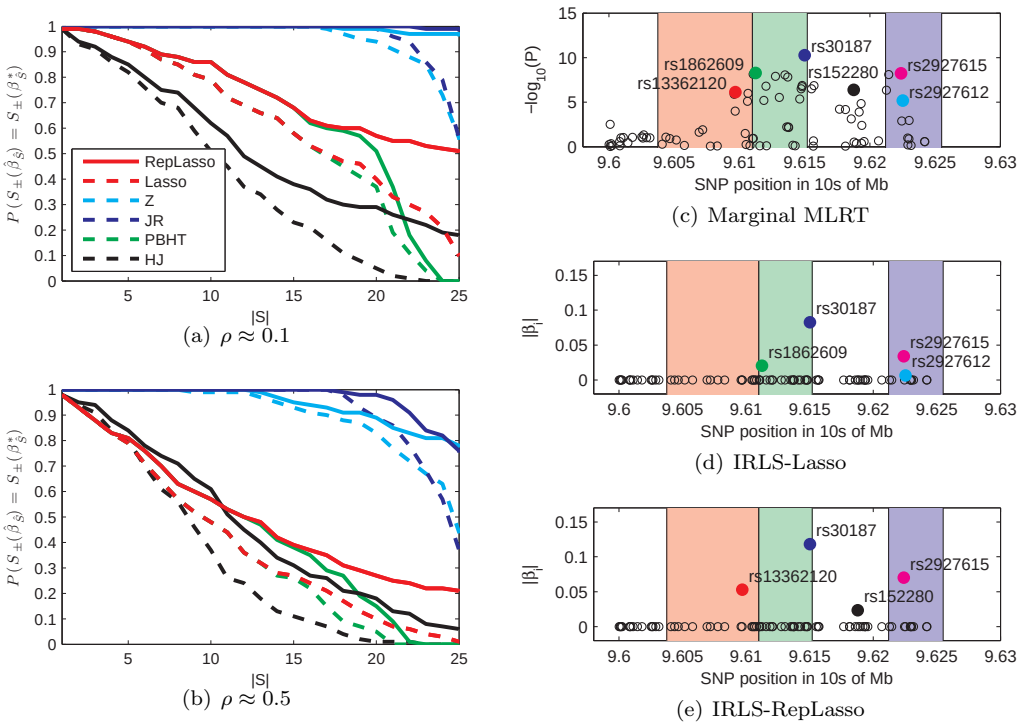

(a) $\rho \approx 0.1$

(b) $\rho \approx 0.5$

(c) Marginal MLRT

(d) IRLS-Lasso

(e) IRLS-RepLasso

Figure 3: (a), (b) Results on synthetic data with $n = 150, p = 50, \mathcal{G}$ = groups of 2 and within-group correlations ρ. We show the empirical probability that a subset of the correct *signed* support is recovered as a function of support size. For each base-method we show two curves, grouped by color. The performance of the existing Lasso variant (i.e., Lasso, Z, JR, PBHT, HJ) is shown as dashed curve; the performance of the algorithm with the Lasso replaced by RepLasso is shown in solid (see text for details). (c) A GWAS Manhattan plot for ankylosing spondylitis on a region of chromosome 5. The x-axis gives SNP position; circles indicate the $-\log_{10}(p)$-values of marginal association tests. The red, green and blue shaded regions indicate the CAST, ERAP1 and ERAP2 genes respectively. Solid circles highlight the SNPs that were chosen by the methods below. (d) The IRLS method of Lee et al. [12] with Lasso. Circles indicate the magnitudes of estimated β coefficients. Solid circles indicate the first 4 SNPs that are chosen by Lasso. The method selects multiple SNPs from the same gene. (e) The RepLasso avoids this if θ is chosen large enough.

A Lars-like Variation. We note at this point that the Lars algorithm [3] is a special case of RepLasso if we set $\theta = \mathbf{0}$ and force $L = 0$. If we only force $L = 0$ but allow $\theta \neq \mathbf{0}$, then the resulting algorithm can be seen as a generalization of Lars. We analyze its performance in the Supplementary Material and show a similar result for its support recovery behavior.

6 Results

Synthetic Data. To underline the findings of Section 5 we focus on a set of experiments which analyses the probability of correctly recovering a subset of the correct *signed* support. We fix $\mathcal{G}, \beta^*$ so that **A1** holds. Conditioned on $\mathcal{G}$ we also sample X with unit-length columns that are independent between groups $G \in \mathcal{G}$ but exhibit some correlation ρ within groups. Given X, β^*, we generate y according to Eq. (1), with $\sigma^2 = 0.2^2$. In Figures 3(a) and 3(b) we investigate the performance of the RepLasso (solid red) and the Lasso (dashed red). Notice that the curve for RepLasso lies above that of Lasso, giving empirical support to Theorems 2 and 3. Additionally,

we evaluate the performance of four other methods that solve a standard Lasso problem after preprocessing the data X, y in some way. For each algorithm we show two curves, grouped by colors: the original method is shown as dashed curve, the method with the Lasso replaced by RepLasso as solid curve. The four methods are: (1) the Adaptive Lasso of Zou [26] (*Z*); (2) the "Whitened" Lasso of Jia and Rohe [9] (*JR*); (3) the Preconditioned Lasso of Paul et al. [15] (*PBHT*); and (4) Correlation Sifting of Huang and Jojic [6] (*HJ*). The results in Figures 3(a) and 3(b) highlight that using RepLasso as a drop-in replacement these algorithms can also be improved. Results are similar when $\mathcal{G}$ contains larger groups or when $n \ll p$.[3]

GWAS application. A second experiment considers the application of the Lasso to a Genome-Wide Association Study (GWAS). A GWAS hopes to find Single Nucleotide Polymorphisms (SNPs) that are associated with disease status. Our focus is on the disease ankylosing spondylitis and a region of chromosome 5, where susceptibility SNPs had been previously reported [1]. The mainstream GWAS methodology tests each form a large set of SNPs *marginally* for association using a maximum likelihood ratio test (MLRT) and plots the resulting p-values on a "Manhattan plot", as in Figure 3(c). Due to linkage disequilibrium, many small p-values lie close to each other. Alternatively, a penalized logistic regression could also be used to regress the SNPs onto disease status, which would then highlight interesting SNPs by the magnitudes of the learned regression coefficients. Lee et al. [12] have proposed an IRLS strategy for estimating an ℓ_1 constrained logistic regression by solving a Lasso problem on a quadratic approximation of the logistic objective. The magnitudes of the first four regression coefficients estimated by this method are shown in Figure 3(d). As can be seen, two pairs of selected SNPs lie near each other in two genes. A researcher might wish to discourage the Lasso from choosing multiple SNPs from the same gene. The RepLasso is ideally suited to this task. Given the gene partition $\mathcal{G}$ (in this case by CAST, ERAP1 and ERAP2 genes) we can replace the Lasso in the IRLS algorithm by the RepLasso and produce a different parameter estimate. If θ is large (e.g., 20 for each group), RepLasso avoids selecting multiple SNPs from the same gene, as seen in Figure 3(e). These SNPs may be worthy of further study.

7 Conclusion

In this paper we presented a homotopy-style algorithm that approximates an underlying nonconvex problem by producing a suitable sequence of surrogates that locally approximate $\Omega_{\theta,\mathcal{G}}(\cdot)$ well. The Lasso approach, in comparison, revolves around finding a single global surrogate that often approximates $\Omega_{\theta,\mathcal{G}}(\cdot)$ poorly. As shown by Theorem 1, our method will in certain cases sweep out a global minima path of $(P1)$. Further, we showed in Section 5 that even though RepLasso may not exactly solve $(P1)$ in general, in relevant regression problems RepLasso will not do worse than the Lasso and in practice often outperforms it.

Several extensions can be considered. Firstly, we defined $\Omega_{\theta,\mathcal{G}}(\cdot)$ as a sum over pairs of variables in groups of the partition $\mathcal{G}$. More flexible constraint functions could potentially be defined if the sum is allowed to be over an arbitrary set of pairs. Secondly, our overall strategy was to decompose the nonconvex constraint balls induced by $\Omega_{\theta,\mathcal{G}}(\cdot)$ as a union of simpler, convex balls. In this work the constraint function $\Omega_{\theta,\mathcal{G}}(\cdot)$ gave rise to a union of weighted ℓ_1 balls. This motivates a more direct definition of nonconvex constraint balls as a union of convex balls. For instance, one could consider unions of weighted ℓ_∞ balls or a mix of weighted ℓ_∞ and weighted ℓ_1 balls. So long as these convex building blocks are consistent with [16] it should still be possible to efficiently compute local minima paths segments as demonstrated in this paper. Thirdly, it would be interesting to see whether results such as in Loh and Wainwright [13] could be extended to argue for statistical consistency of the RepLasso in cases where local minima paths are produced. An important ingredient of such an analysis will be that $\Omega_{\theta,\mathcal{G}}(\beta)$ is not "too" nonconvex, which might hold if θ is sufficiently small.

[3] If $n < p$ we let the Adaptive Lasso scale columns of X by univariate regression coefficients.

Acknowledgments

We thank Francis Bach and Alexander Young for helpful comments and Nebojsa Jojic for early input out of which this research grew.

References

[1] P.R. Burton et al. Association scan of 14,500 nonsynonymous SNPs in four diseases identifies autoimmunity variants. *Nature genetics*, 39(11):1329–1337, 2007.

[2] E.J. Candes, M.B. Wakin, and S.P. Boyd. Enhancing sparsity by reweighted ℓ_1 minimization. *J. Fourier Anal. Appl.*, 14(5-6):877–905, 2008.

[3] B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least Angle Regression. *Ann. Stat.*, 32:407–499, 2004.

[4] J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *J. Amer. Statist. Assoc.*, 96:1348–1360, 2001.

[5] T. Hastie, S. Rosset, R. Tibshirani, and J. Zhu. The entire regularization path for the Support Vector Machine. *J. Mach. Learn. Res.*, 5:1391–1415, December 2004.

[6] J.C. Huang and N. Jojic. Variable selection through Correlation Sifting. In V. Bafna and S.C. Sahinalp, editors, *RECOMB*, volume 6577 of *LNCS*, pages 106–123. Springer, 2011.

[7] D.R. Hunter and R. Li. Variable selection using MM algorithms. *Ann. Stat.*, 33(4):1617, 2005.

[8] R. Jenatton, J.-Y. Audibert, and F. Bach. Structured variable selection with sparsity-inducing norms. *J. Mach. Learn. Res.*, 12:2777–2824, 2011.

[9] J. Jia and K. Rohe. "Preconditioning" to comply with the irrepresentable condition. 2012.

[10] S. Kim and E. P. Xing. Tree-guided Group Lasso for multi-response regression with structured sparsity, with applications to eQTL mapping. *Ann. Appl. Stat.*, 2012.

[11] M. Kowalski and B. Torrésani. Sparsity and persistence: mixed norms provide simple signal models with dependent coefficients. *Signal, Image and Video processing*, 3(3):251–264, 2009.

[12] S.-I. Lee, H. Lee, P. Abbeel, and A.Y. Ng. Efficient ℓ_1 regularized logistic regression. In *Proc. Conf. AAAI Artif. Intell.*, volume 21, page 401, 2006.

[13] P.-L. Loh and M.J. Wainwright. Regularized M-estimators with nonconvexity: Statistical and algorithmic theory for local optima. In *Adv. Neur. Inf. Process. Syst. 26*, pages 476–484. 2013.

[14] M.R. Osborne, B. Presnell, and B.A. Turlach. A new approach to variable selection in least squares problems. *IMA J. Numer. Anal.*, 20(3):389–403, 2000.

[15] D. Paul, E. Bair, T. Hastie, and R. Tibshirani. "Preconditioning" for feature selection and regression in high-dimensional problems. *Ann. Stat.*, 36(4):1595–1618, 2008.

[16] S. Rosset and J. Zhu. Piecewise linear regularized solution paths. *Ann. Stat.*, pages 1012–1030, 2007.

[17] R. Tibshirani. Regression shrinkage and selection via the Lasso. *J. R. Stat. Soc. Ser. B*, 58(1):267–288, 1994.

[18] R.J. Tibshirani. The Lasso problem and uniqueness. *Electronic Journal of Statistics*, 7:1456–1490, 2013.

[19] M.J. Wainwright. Sharp thresholds for high-dimensional and noisy sparsity recovery using ℓ_1-constrained quadratic programming (Lasso). *IEEE Trans. Inf. Theo.*, 55(5):2183–2202, 2009.

[20] L. Wang, G. Chen, and H. Li. Group SCAD regression analysis for microarray time course gene expression data. *Bioinformatics*, 23(12):1486–1494, 2007.

[21] F.L. Wauthier, N. Jojic, and M.I. Jordan. A comparative framework for preconditioned Lasso algorithms. In *Adv. Neural Inf. Process. Syst. 26*, pages 1061–1069. 2013.

[22] M. Yuan and Y. Lin. Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc., Ser. B*, 68(1):49–67, 2006.

[23] C.-H. Zhang. Nearly unbiased variable selection under Minimax Concave Penalty. *Ann. Stat.*, 38(2):894–942, 2010.

[24] P. Zhao and B. Yu. On model selection consistency of Lasso. *J. Mach. Learn. Res.*, 7:2541–2563, 2006.

[25] Y. Zhou, R. Jin, and S. Hoi. Exclusive Lasso for multi-task feature selection. In *Artificial Intelligence and Statistics*, pages 988–995, 2010.

[26] H. Zou. The adaptive Lasso and its oracle properties. *J. Amer. Statist. Assoc.*, 101(476):1418–1429, 2006.

[27] H. Zou and T. Hastie. Regularization and variable selection via the Elastic Net. *J. R. Stat. Soc., Ser. B*, 67:301–320, 2005.

[28] H. Zou and R. Li. One-step sparse estimates in nonconcave penalized likelihood models. *Ann. Stat.*, 36(4):1509, 2008.

A Proofs of Section 3

Recall that given a partition $\mathcal{G} = \{G_1, \ldots, G_g\}$ of $\{1, \ldots, p\}$ without singleton or empty sets, and a vector $\theta = (\theta_1, \ldots, \theta_g) \geq 0$, we defined

$$\Omega_{\theta,\mathcal{G}}(\beta) = \sum_{i<j \in G_{g'} \in \mathcal{G}} \frac{\omega_{\theta_{g'}}(\beta_i, \beta_j)}{|G_{g'}| - 1} \qquad \omega_{\theta_{g'}}(\beta_i, \beta_j) = \min(|\beta_i|, |\beta_j|)(1 + \theta_{g'}) + \max(|\beta_i|, |\beta_j|). \tag{8}$$

Let $B_{\theta,\mathcal{G}}(\tau) = \{\beta \in \mathbb{R}^p : \Omega_{\theta,\mathcal{G}}(\beta) \leq \tau\}$ be the induced constraint balls and $\Gamma(i) \in \{1, \ldots, g\}$ the (unique) group so that $i \in G_{\Gamma(i)}$.

A.1 Proof of Proposition 1

Proposition 1 (Union Decomposition)**.** *Let the partition be $\mathcal{G} = \{G_1, \ldots, G_g\}$ and the parameter $\theta = (\theta_1, \ldots, \theta_g) \geq 0$. There is a finite set $\mathcal{S}_{\theta,\mathcal{G}} \subset \mathbb{R}^p$ of vectors $s \geq \mathbf{1}$, so that for any $\tau > 0$*

$$B_{\theta,\mathcal{G}}(\tau) = \bigcup_{s \in \mathcal{S}_{\theta,\mathcal{G}}} B_s(\tau). \tag{9}$$

Define $\Pi_{g'}$ to be all permutations $\pi_{g'}$ of the elements in $G_{g'}$ and let $\Pi_{\mathcal{G}} = \times_{g'=1}^{g} \Pi_{g'}$ be their cross-product, whose elements $\pi \in \Pi_{\mathcal{G}}$ are g-tuples of permutations $\pi = (\pi_1, \ldots, \pi_g)$. For some $\pi \in \Pi_{\mathcal{G}}$, denote by $\pi_{\Gamma(i)}(i) \in \{1, \ldots, |G_{\Gamma(i)}|\}$ the position of $i \in G_{\Gamma(i)}$ in permutation $\pi_{\Gamma(i)}$. We have

$$\mathcal{S}_{\theta,\mathcal{G}} = \cup_{\pi \in \Pi_{\mathcal{G}}} \{s_\pi\} \tag{10}$$

$$s_{\pi,i} = 1 + (\pi_{\Gamma(i)}(i) - 1) \frac{\theta_{\Gamma(i)}}{|G_{\Gamma(i)}| - 1} \ \forall i = 1, \ldots, p. \tag{11}$$

Proof. We first show $B_{\theta,\mathcal{G}}(\tau) \subseteq \bigcup_{s \in \mathcal{S}_{\theta,\mathcal{G}}} B_s(\tau)$. Consider some $\beta \in B_{\theta,\mathcal{G}}(\tau)$ and let $\pi = (\pi_1, \ldots, \pi_g)$ be a tuple of permutations (not necessarily unique) induced by sorting the elements $|\beta_i|$ within each group specified by $\mathcal{G}$ so that for each group index g' we have

$$|\beta_{\pi_{g'}^{-1}(1)}| \geq \ldots \geq |\beta_{\pi_{g'}^{-1}(|G_{g'}|)}|. \tag{12}$$

By construction of $\Omega_{\theta,\mathcal{G}}(\cdot)$, β lies in the set

$$\left\{ \beta' \in \mathbb{R}^p : \sum_{g'=1}^{g} \sum_{i=1}^{|G_{g'}|} \left(\frac{(|G_{g'}| - i) + (i-1)(1 + \theta_{g'})}{|G_{g'}| - 1} \right) |\beta'_{\pi_{g'}^{-1}(i)}| \leq \tau \right\} \tag{13}$$

$$= \left\{ \beta' \in \mathbb{R}^p : \sum_{g'=1}^{g} \sum_{i=1}^{|G_{g'}|} \left(1 + \frac{(i-1)\theta_{g'}}{|G_{g'}| - 1} \right) |\beta'_{\pi_{g'}^{-1}(i)}| \leq \tau \right\} \tag{14}$$

$$= \left\{ \beta' \in \mathbb{R}^p : \sum_{g'=1}^{g} \sum_{i \in G_{g'}} \left(1 + (\pi_{g'}(i) - 1) \frac{\theta_{g'}}{|G_{g'}| - 1} \right) |\beta'_i| \leq \tau \right\} \tag{15}$$

$$= \left\{ \beta' \in \mathbb{R}^p : \sum_{i=1}^{p} \left(1 + \left(\pi_{\Gamma(i)}(i) - 1\right) \frac{\theta_{\Gamma(i)}}{|G_{\Gamma(i)}| - 1} \right) |\beta'_i| \leq \tau \right\} = B_{s_\pi}(\tau) \tag{16}$$

We therefore conclude that $B_{\theta,\mathcal{G}}(\tau) \subseteq \bigcup_{s \in \mathcal{S}_{\theta,\mathcal{G}}} B_s(\tau)$, with

$$\mathcal{S}_{\theta,\mathcal{G}} = \cup_{\pi \in \Pi_{\mathcal{G}}} \{s_\pi\} \tag{17}$$

$$s_{\pi,i} = 1 + (\pi_{\Gamma(i)}(i) - 1) \frac{\theta_{\Gamma(i)}}{|G_{\Gamma(i)}| - 1} \ \forall i = 1, \ldots, p. \tag{18}$$

For the other direction, suppose that $\beta \in B_{s_{\tilde{\pi}}}(\tau)$ for some arbitrary tuple of permutations $\tilde{\pi} \in \Pi_{\mathcal{G}}$, which means that

$$\sum_{i=1}^{p} \left(1 + \left(\tilde{\pi}_{\Gamma(i)}(i) - 1\right) \frac{\theta_{\Gamma(i)}}{|G_{\Gamma(i)}| - 1}\right) |\beta_i| \leq \tau. \tag{19}$$

Then notice that if π is a (not necessarily unique) tuple of permutations induced by ordering elements $|\beta_i|$ within groups, we have, by arguing from pairwise swaps within groups that take $\tilde{\pi}$ to π, that

$$\sum_{i=1}^{p} \left(1 + \left(\pi_{\Gamma(i)}(i) - 1\right) \frac{\theta_{\Gamma(i)}}{|G_{\Gamma(i)}| - 1}\right) |\beta_i| \leq \sum_{i=1}^{p} \left(1 + \left(\tilde{\pi}_{\Gamma(i)}(i) - 1\right) \frac{\theta_{\Gamma(i)}}{|G_{\Gamma(i)}| - 1}\right) |\beta_i| \leq \tau, \tag{20}$$

and so $\beta \in B_{\theta,\mathcal{G}}(\tau)$. It follows that $\bigcup_{s \in \mathcal{S}_{\theta,\mathcal{G}}} B_s(\tau) \subseteq B_{\theta,\mathcal{G}}(\tau)$ and so $B_{\theta,\mathcal{G}}(\tau) = \bigcup_{s \in \mathcal{S}_{\theta,\mathcal{G}}} B_s(\tau)$.

□

A.2 Proof of Proposition 2

Recall that we are considering the nonconvex optimization problem

$$\begin{aligned}\beta(\tau) &\in \operatorname*{argmin}_{\beta\in\mathbb{R}^p} J_\tau(\beta) \\ &= \operatorname*{argmin}_{\beta\in\mathbb{R}^p} \begin{cases} \frac{1}{2n}\|y - X\beta\|_2^2 & \text{if } \beta \in B_{\theta,\mathcal{G}}(\tau) \\ \infty & \text{o.w.} \end{cases}.\end{aligned} \tag{P1}$$

Proposition 2 (Local Piecewise Linearity)**.** *Suppose X has absolutely continuous distribution and that $\exists \tau' > 0$ s.t. $\nexists \beta \in B_{\theta,\mathcal{G}}(\tau')$ which is a minimum of $\|y - X\beta\|_2^2$. Let $\tau_{\max}$ be the supremum over these τ'. The set of local minima of $J_\tau(\cdot)$ in (P1) with $\tau \in (0, \tau_{\max})$ is w.p. 1 a finite union of piecewise linear paths, each path indexed by τ and lying on a ball $B_s(\tau), s \in \mathcal{S}_{\theta,\mathcal{G}}$.*

Proof. Given the assumptions, for all $\tau \in (0, \tau_{\max})$, the elements β on the boundary of $B_{\theta,\mathcal{G}}(\tau)$ satisfy $(y - X\beta)^\top X \neq 0$. For each $\tau \in (0, \tau_{\max})$, let $\mathcal{M}_{\theta,\mathcal{G}}(\tau)$ be the set of local minima of $J_\tau(\cdot)$. Let the set $\mathcal{S}_{\theta,\mathcal{G}}$ be defined as in Proposition 1: For $\Pi_{\mathcal{G}}$ the set of g-tuples of permutations induced by $\mathcal{G}$,

$$\mathcal{S}_{\theta,\mathcal{G}} = \cup_{\pi\in\Pi_{\mathcal{G}}}\{s_\pi\} \tag{21}$$

$$s_{\pi,i} = 1 + (\pi_{\Gamma(i)}(i) - 1)\frac{\theta_{\Gamma(i)}}{|G_{\Gamma(i)}| - 1} \ \forall i = 1, \ldots, p. \tag{22}$$

For some $s_\pi \in \mathcal{S}_{\theta,\mathcal{G}}$, define $\mathcal{M}_{s_\pi}(\tau)$ to be the solution to $(P1)$ with $B_{\theta,\mathcal{G}}(\tau)$ replaced by $B_{s_\pi}(\tau)$. For each $s_\pi \in \mathcal{S}_{\theta,\mathcal{G}}$ the ball $B_{s_\pi}(\tau)$ corresponds to a weighted ℓ_1 norm, and if X is drawn from an absolutely continuous distribution, then the solution $\mathcal{M}_{s_\pi}(\tau)$ is with probability 1 unique on $(0, \tau_{\max})$ [18]. Additionally, the result of Rosset and Zhu [16] shows that the resulting regularization path $\mathcal{M}_{s_\pi}(\tau)$ is piecewise linear on $(0, \tau_{\max})$. Due to the union decomposition of Proposition 1, it follows immediately that $\mathcal{M}_{\theta,\mathcal{G}}(\tau) \subseteq \bigcup_{s_\pi\in\mathcal{S}_{\theta,\mathcal{G}}} \mathcal{M}_{s_\pi}(\tau)$ for $\tau \in (0, \tau_{\max})$. However, we seek not a superset of $\mathcal{M}_{\theta,\mathcal{G}}(\tau)$, but a characterisation as a union of paths on the boundaries of weighted ℓ_1 balls. That, is we seek a set $P \subseteq \Pi_{\mathcal{G}}$ so that $\mathcal{M}_{\theta,\mathcal{G}}(\tau) = \bigcup_{\pi\in P} \mathcal{M}_{s_\pi}(\tau)$ for $\tau \in (0, \tau_{\max})$. The existence of such a set P can be guaranteed if for any $s_\pi \in \mathcal{S}_{\theta,\mathcal{G}}$, $\mathcal{M}_{s_\pi}(\tau)$ either lies $\forall \tau \in (0, \tau_{\max})$ in the interior of $B_{\theta,\mathcal{G}}(\tau)$ or it lies $\forall \tau \in (0, \tau_{\max})$ on the boundary of $B_{\theta,\mathcal{G}}(\tau)$. To show this, we show that for $\tau \in (0, \tau_{\max})$ no local minimum in $\mathcal{M}_{\theta,\mathcal{G}}(\tau)$ lies at a concave kink of $B_{\theta,\mathcal{G}}(\tau)$ (which are the points where a path would switch from being in the interior to being on the boundary or vice versa).

Suppose then (for the purpose of deriving a contradiction) that for some $\tau \in (0, \tau_{\max})$, we have that β is a local minimum in $\mathcal{M}_{\theta,\mathcal{G}}(\tau)$ that lies at one of the concave kinks of $B_{\theta,\mathcal{G}}(\tau)$. If β lies at a concave kink, then since $\tau_{\max} > 0$, we know that for at least two elements $i \neq j \in G \in \mathcal{G}$, $\beta_i \neq 0, \beta_j \neq 0$. For if only a single element $\neq 0$, then we lie at one of the points of $B_{\theta,\mathcal{G}}(\tau)$ and if the only two nonzero elements lie in different groups, β cannot lie at a concave kink. Specifically, the concave kink is identified by sets of indices i in a group $G \in \mathcal{G}$ so that the corresponding $\beta_i \neq 0$ have identical magnitude. The vector β induces a set $\Sigma \subseteq \Pi_{\mathcal{G}}$ of g-tuples of permutations σ by sorting $|\beta_i|$ by their magnitudes within each group $G \in \mathcal{G} = \{G_1, \ldots, G_g\}$ (with tie-breaking). We know that for each $\sigma \in \Sigma$, $\|\text{diag}(s_\sigma)\beta\|_1 = \tau$, that is, β lies on the boundary of $B_{s_\sigma}(\tau)$. Each σ thus corresponds to an active constraint on β. Since we can think of β as a local minimum of $\|y - X\beta\|_2^2$, subject to either of these (convex) constraints, we have by convexity for any $\sigma \in \Sigma$ a subgradient vector $z_\sigma \in \partial\|\beta\|_1$ and a constant λ_σ so that

$$(y - X\beta)^\top X = \lambda_\sigma \text{diag}(z_\sigma)s_\sigma. \tag{23}$$

Because there are at least two elements $i \neq j \in G \subseteq \mathcal{G}$ with $|\beta_i| = |\beta_j| \neq 0$ we know that $\forall \sigma \in \Sigma$, $z_{\sigma,i} = \text{sgn}(\beta_i), z_{\sigma,j} = \text{sgn}(\beta_j)$, which implies that $z_{\sigma,i}s_{\sigma,i} \neq 0, z_{\sigma,j}s_{\sigma,j} \neq 0$. Additionally, by construction $(y - X\beta)^\top X \neq 0$ and so we know $\lambda_\sigma \neq 0$. By the construction of s_π in Eq. (22),

we know that $\exists \sigma_1 \neq \sigma_2 \in \Sigma$, so that s_{σ_1} and s_{σ_2} differ only on elements i, j. However Eq. (23) then cannot simultaneously hold unless $\lambda_\sigma = 0$ and $(y - X\beta)^\top X = 0$ which we ruled out earlier. Thus we have a contradiction and so the assumption that β lies at a concave kink must be wrong.

Because local minima in $\mathcal{M}_{\theta,\mathcal{G}}(\tau)$ never lie at concave kinks of $B_{\theta,\mathcal{G}}(\tau)$ for $\tau \in (0, \tau_{\max})$, we know that for each local minimum path on $(0, \tau_{\max})$, there is a $\pi \in \Pi_{\mathcal{G}}$ so that the path lies on $B_{s_\pi}(\tau)$. That is, there is some nonempty subset $P \subseteq \Pi_{\mathcal{G}}$ so that $\mathcal{M}_{\theta,\mathcal{G}}(\tau) = \bigcup_{\pi \in P} \mathcal{M}_{s_\pi}(\tau)$ is a union of piecewise linear paths. $\square$

B Proofs of Section 4

B.1 Proof of Proposition 3

Recall that we are considering the surrogate problem

$$\bar{\beta}(\lambda) \in \operatorname{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda \|\operatorname{diag}(s^*)\beta\|_1 . \tag{S}$$

For a positive vector b, let $\bar{\beta}_b(\lambda)$ be a solution to (S) with penalty $\lambda\|\operatorname{diag}(b)\beta\|_1$.

Proposition 3 (Recoverability of (S))**.** *Suppose X has absolutely continuous distribution. For any vectors $a \geq b \geq \mathbf{1}$ and $\lambda > 0$, w.p. 1 $\bar{\beta}_a(\lambda), \bar{\beta}_b(\lambda)$ are unique. If additionally $\|diag(a)\bar{\beta}_b(\lambda)\|_1 = \|diag(b)\bar{\beta}_b(\lambda)\|_1$, then $\bar{\beta}_a(\lambda) = \bar{\beta}_b(\lambda)$.*

Proof. Since X is absolutely continuous, $a \geq b \geq 0$ and $\lambda > 0$, it follows that $\bar{\beta}_a(\lambda), \bar{\beta}_b(\lambda)$ are almost surely unique [18]. Since $a \geq b \geq \mathbf{1}$, we have $\forall \beta \in \mathbb{R}^p$

$$\frac{1}{2n} \|y - X\beta\|_2^2 + \lambda \|\operatorname{diag}(b)\beta\|_1 \leq \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda \|\operatorname{diag}(a)\beta\|_1 . \tag{24}$$

However, we also know

$$\lambda \left\|\operatorname{diag}(b)\bar{\beta}_b(\lambda)\right\|_1 = \lambda \left\|\operatorname{diag}(a)\bar{\beta}_b(\lambda)\right\|_1 . \tag{25}$$

It follows that we must have $\bar{\beta}_a(\lambda) = \bar{\beta}_b(\lambda)$. □

C Proofs of Section 5

Section 5 compares the estimator of β^* produced by the RepLasso algorithm, with the estimator of β^* produced by the Lasso.

The RepLasso is a generalization of the Lasso homotopy method, which maintains a set of weights $s(\lambda)$. Indeed, the RepLasso is identical to the Lars algorithm with Lasso modification of Efron et al. [3] if we force $\theta = \mathbf{0}$, which implies that $\forall\lambda, s(\lambda) = \mathbf{1}$ (We note, however, that for notational convenience the definition of $\bar{w}_A$ differs slightly from that in Efron et al. [3] in that case). In the following we will carry out our comparison of RepLasso with the Lasso homotopy method by comparing the RepLasso with $\theta \neq \mathbf{0}$ and the RepLasso with $\theta = \mathbf{0}$. We will denote by $\hat{\beta}(\lambda)$ the estimator resulting from the specialization to the Lasso case. Similarly, we let $\hat{w}_A$ be the vector corresponding to $\bar{w}_A$ for the Lasso specialization.

We use the following notation inspired by Wainwright [19] and Wauthier et al. [21]. Suppose that the support set of β^* is $S \triangleq S(\beta^*)$. Let X_j be the column j of X and X_A a matrix which consists of the columns indexed by A. For all $j \in S^c$ and $i \in S$, let

$$\mu_j = X_j^\top X_S (X_S^\top X_S)^{-1} \mathrm{sgn}(\beta_S^*) \qquad \eta_j = X_j^\top \left(I_{n\times n} - X_S (X_S^\top X_S)^{-1} X_S^\top\right) \frac{w}{n} \tag{26}$$

$$\gamma_i = e_i^\top \left(\frac{1}{n} X_S^\top X_S\right)^{-1} \mathrm{sgn}(\beta_S^*) \qquad \epsilon_i = e_i^\top \left(\frac{1}{n} X_S^\top X_S\right)^{-1} X_S^\top \frac{w}{n}. \tag{27}$$

The proofs of Section 5 use subsets of the following assumptions.

A1: $\forall G \in \mathcal{G}, |\{i \in G : \beta_i^* \neq 0\}| \leq 1$

A2: $\forall A \subset S$ and u_A the equiangular vector in Eq. (2.6) of [3], $\nexists j \in A^c, |X_A^\top u_A| = |X_j^\top u_A|\mathbf{1}$

A3: $X_S^\top X_S$ is invertible

A4: $|\mu_j| < 1, \forall j \in S^c$, $\mathrm{sgn}(\beta_i^*)\gamma_i > 0 \ \ \forall i \in S$.

C.1 Proof of Theorem 2

Theorem 2 (Support Subset Recovery)**.** *Assume that* **A1–2** *hold. Denote by* $\hat{\beta}(\lambda)$ *and* $\bar{\beta}(\lambda)$ *the Lasso and RepLasso solutions using penalty parameter* λ*. Conditioned on* X, y*, we have for any* $\lambda_{\min} > 0$

$$\forall \lambda \geq \lambda_{\min}\ S(\hat{\beta}(\lambda)) \subseteq S \implies \forall \lambda \geq \lambda_{\min}\ S(\bar{\beta}(\lambda)) \subseteq S.$$

Proof. Suppose then that $\forall \lambda \geq \lambda_{\min}, S(\hat{\beta}(\lambda)) \subseteq S$. Suppose that $\hat{\beta}(\lambda_{\min})$ corresponds to iteration t of the Lasso. For $t' \leq t$, let $\hat{A}_{t'}$ and $\bar{A}_{t'}$ be the sequence of active sets of the Lasso and RepLasso up to iteration t. With a slight abuse of notation we will temporarily treat an active set as an unordered set. Assumption **A2** guarantees that for the Lasso, any variable that is at some point in the active set is also at some point in the support set. To see this, note that that by **A2**, the vector $\hat{w}_A$ never contains a zero element. If it did, then an equiangular vector u_A of X_A as in Eq. (2.6) of [3] could be constructed using a strict subset of vectors indexed by A, violating assumption **A2**. But if $\hat{w}_A$ does not contain a zero element, then the elements in the active set A cannot indefinitely be assigned a $\hat{\beta}_A(\lambda)$ coefficient of zero as λ is swept out. Finally, because we know $\forall \lambda \geq \lambda_{\min}, S(\hat{\beta}(\lambda)) \subseteq S$, this means that $\forall t' \leq t, \hat{A}_{t'} \subseteq S$. We will now argue by induction that the induced sequence of active sets $\bar{A}_{t'}$ of the RepLasso also satisfies $\forall t' \leq t, \bar{A}_{t'} \subseteq S$.

Base case: Since $s(\|X^\top y\|_\infty) = \mathbf{1}$, the first variable selected by RepLasso and the Lasso method is the same. That is, $\hat{A}_1 = \bar{A}_1 \subseteq S$ at iteration 1.

Inductive step: Assume that $\forall t'' \leq t', \hat{A}_{t''} = \bar{A}_{t''} \subseteq S$. Since $\forall t'' \leq t', \bar{A}_{t''} \subseteq S$, we know by **A1** that for all λ up to iteration t', $s(\lambda)$ did not change on S, i.e. $s_S(\lambda) = \mathbf{1}$. This in particular means that both the Lasso and the RepLasso will have arrived at the same value of λ and intermediate estimate $\hat{\beta}(\lambda) = \bar{\beta}(\lambda)$ of β^* at the end of stage 4 of iteration $t' - 1$ and the same vectors $\hat{w}_A = \bar{w}_A$ at the end of stage 2 of iteration t'. To see this, notice that since $\forall t'' \leq t', \bar{A}_{t''} \subseteq S$, and since for all λ up to iteration t' we had $s_S(\lambda) = \mathbf{1}$, the RepLasso is up to stage 2 of iteration t' equivalent to running Lasso on the subset of variables X_S, y.

At stage 3 of iteration t', the RepLasso algorithm determines whether to add or remove a variable from $\bar{A}_{t'}$ in stage 1 of iteration $t' + 1$. Since the value of λ and the intermediate variables $\hat{\beta}(\lambda) = \bar{\beta}(\lambda)$ and $\hat{w}_A = \bar{w}_A$ are the same at stage 2 of iteration t' we can now use properties of $s(\lambda)$ to show that this implies $\hat{A}_{t'+1} = \bar{A}_{t'+1}$. We consider two cases:

1. The Lasso determines to add a variable in iteration $t' + 1$ (first bullet in stage 3). Since $s_S(\lambda) = \mathbf{1}$ did not change on S, since we always have $s(\lambda) \geq \mathbf{1}$ and since $\hat{A}_{t'+1} \subseteq S$, it follows that the RepLasso will add the same variable.

2. The Lasso determines to remove a variable in iteration $t'+1$ (second bullet in stage 3). Since $s_S(\lambda) = \mathbf{1}$ did not change on S and since we always have $s(\lambda) \geq \mathbf{1}$, it then follows that the RepLasso will remove the same variable.

Hence, it follows that $\hat{A}_{t'+1} = \bar{A}_{t'+1} \subseteq S$. By the principle of induction, we have shown that $\forall t' \leq t, \bar{A}_{t'} \subseteq S$.

As the value of λ at the end of stage 4 of iteration t must be the same for RepLasso and Lasso, and since for that value we have by definition $\lambda < \lambda_{\min}$, we now know that $\forall \lambda \geq \lambda_{\min}, S(\bar{\beta}(\lambda)) \subseteq S$. □

C.2 Proof of Theorem 3

Theorem 3 (Signed Support Recovery)**.** *Assume that* **A1**–**4** *hold. Denote by* $\hat{\beta}(\lambda)$ *and* $\bar{\beta}(\lambda)$ *the Lasso and RepLasso solutions using penalty parameter* λ*. For any* $\lambda_{\min} > 0$*, we have with probability 1 over an absolutely continuous distribution on noise* w

$$\forall \lambda \geq \lambda_{\min}\ S(\hat{\beta}(\lambda)) \subseteq S, S_{\pm}(\hat{\beta}(\lambda_{\min})) = S_{\pm} \implies \forall \lambda \geq \lambda_{\min}\ S(\bar{\beta}(\lambda)) \subseteq S, S_{\pm}(\bar{\beta}(\lambda_{\min})) = S_{\pm}.$$

Proof. Suppose then that $\forall \lambda \geq \lambda_{\min}\ S(\hat{\beta}(\lambda)) \subseteq S, S_{\pm}(\hat{\beta}(\lambda_{\min})) = S_{\pm}$. Suppose that $\hat{\beta}(\lambda_{\min})$ corresponds to iteration t of the Lasso and for $t' \leq t$ define $\hat{A}_{t'}$ and $\bar{A}_{t'}$ to be the active set of Lasso and RepLasso at iteration t'. With a slight abuse of notation we will temporarily treat an active set as an unordered set. Using the same reasoning as in Theorem 2, it follows from $\forall \lambda \geq \lambda_{\min}\ S(\hat{\beta}(\lambda)) \subseteq S$ and **A2** that $\forall t' \leq t\ \hat{A}_{t'} \subseteq S$ and from this via **A1** that $\forall t' \leq t, \bar{A}_{t'} \subseteq S$. The latter implies that $\forall \lambda \geq \lambda_{\min}, S(\bar{\beta}(\lambda)) \subseteq S$. It remains to be shown that $S_{\pm}(\bar{\beta}(\lambda_{\min})) = S_{\pm}$.

By **A3**–**4** we assumed that $X_S^\top X_S$ is invertible, $|\mu_j| < 1, \forall j \in S^c$ and $\operatorname{sgn}(\beta_i^*)\gamma_i > 0, \forall i \in S$. We can thus apply Lemma 1 of Wauthier et al. [21], which holds with probability 1 over noise instances w. We thus know that with probability 1, $S_{\pm}(\hat{\beta}(\lambda_{\min})) = S_{\pm} \iff \lambda_l < \lambda_{\min} < \lambda_u$, where

$$\lambda_l = \max_{j \in S^c} \frac{\eta_j}{(2[\![\eta_j > 0]\!] - 1) - \mu_j} \qquad \lambda_u = \min_{i \in S} \left| \frac{\beta_i^* + \epsilon_i}{\gamma_i} \right|_+, \tag{28}$$

$[\![\cdot]\!]$ denotes the indicator function and $\eta_j, \mu_j, \epsilon_i, \gamma_i$ are defined on X, β^*, w as in Eqs. (26, 27). The RepLasso algorithm traces out the solution path of a sequence of weighted Lasso problems parameterized by $s(\lambda)$

$$\bar{\beta}(\lambda) \in \operatorname{argmin}_{\beta \in \mathbb{R}^p} \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda \|\operatorname{diag}(s(\lambda))\beta\|_1. \tag{29}$$

Suppose we temporarily decouple $s(\lambda)$ from the penalty parameter λ and fix it at $s(\lambda_{\min})$. We will show that we can then with probability 1 apply Lemma 1 of Wauthier et al. [21] to the resulting $s(\lambda_{\min})$-weighted Lasso problem by applying it to the unweighted Lasso problem on $\bar{X} = X\operatorname{diag}(s(\lambda_{\min}))^{-1}, \bar{\beta}^* = \operatorname{diag}(s(\lambda_{\min}))\beta^*, \bar{w} = w$. Let $\bar{\eta}_j, \bar{\mu}_j, \bar{\epsilon}_i, \bar{\gamma}_i$ be the corresponding variables defined for $\bar{X}, \bar{\beta}^*, \bar{w}$. We always have $s(\lambda_{\min}) \geq \mathbf{1}$, and since $\forall t' \leq t, \bar{A}_{t'} \subseteq S$, we also know by **A1** that $\forall i \in S, s_i(\lambda_{\min}) = 1$. It follows that $\bar{\eta}_j = \eta_j / s_j(\lambda_{\min})$, $\bar{\mu}_j = \mu_j / s_j(\lambda_{\min})$, $\bar{\epsilon}_i = \epsilon_i$, $\bar{\gamma}_i = \gamma_i$ and $\bar{\beta}_i^* = \beta_i^*$. By **A3**–**4** we assumed that $X_S^\top X_S$ is invertible, $|\mu_j| < 1, \forall j \in S^c$ and $\operatorname{sgn}(\beta_i^*)\gamma_i > 0, \forall i \in S$. Since $s(\lambda_{\min}) \geq \mathbf{1}$ we have that $\bar{X}_S^\top \bar{X}_S$ is invertible, $|\bar{\mu}_j| < 1, \forall j \in S^c$ and $\operatorname{sgn}(\bar{\beta}_i^*)\bar{\gamma}_i > 0, \forall i \in S$. Hence, we are licensed to apply Lemma 1 of Wauthier et al. [21] to the new problem instance which produces new bounds $\bar{\lambda}_l, \bar{\lambda}_u$ in terms of $\bar{\eta}_j, \bar{\mu}_j, \bar{\epsilon}_i, \bar{\gamma}_i$. Because $s(\lambda_{\min}) \geq \mathbf{1}$, simple calculations show that $\bar{\lambda}_l \leq \lambda_l$ and $\lambda_u = \bar{\lambda}_u$. Since the Lasso has (with probability 1) $\lambda_l < \lambda_{\min} < \lambda_u$, the same $\lambda_{\min}$ thus also satisfies $\bar{\lambda}_l \leq \lambda_l < \lambda_{\min} < \lambda_u = \bar{\lambda}_u$. Combining the latter fact with Lemma 1 of Wauthier et al. [21] we then see that with probability 1 $S_{\pm}(\bar{\beta}(\lambda_{\min})) = S_{\pm}$. □

D RepLars: A RepLasso variant

As noted in the paper, if $\theta = \mathbf{0}$, and we force $L = 0$ then the RepLasso algorithm reduces to the Lars algorithm of Efron et al. [3]. (We note, however, that the definition of $\bar{w}_A$ differs slightly from that of the Lars algorithm given in [3]). When we force $L = 0$ but allow $\theta \neq \mathbf{0}$ we have a new algorithm, which we call RepLars. In this section we present this algorithm and analyze its behavior.

Algorithm 2: REPLARS$(X, y, \mathcal{G}, \theta)$

$\bar{y} = 0, A = (), \lambda = \|X^\top y\|_\infty, s(\lambda) = \mathbf{1}, \bar{\beta}(\lambda) = 0$
while $\lambda > 0$

Stage 1
$$\begin{cases} A = (A, i^*), \text{ where } i^* = \text{argmax}_{j \in A^c} \left| X_j^\top (y - \bar{y})/s_j(\lambda) \right| \\ s_M(\lambda^-) = s_M(\lambda) + \frac{\theta_{\Gamma(i^*)}}{|G_{\Gamma(i^*)}|-1}\mathbf{1}, \quad s_{M^c}(\lambda^-) = s_{M^c}(\lambda), \text{ with } M = \left\{A^c \cap G_{\Gamma(i^*)}\right\} \end{cases}$$

Stage 2
$$\left\{ \bar{w}_A = A_A \left(X_A^\top X_A\right)^{-1} \text{diag}(\text{sgn}\left(X_A^\top (y - \bar{y})\right)) s_A(\lambda), \text{ with } A_A \text{ s.t. } \|X_A \bar{w}_A\|_2^2 = 1 \right.$$

Stage 3
$$\left\{ \text{Find smallest } \rho > 0 \text{ s.t. } \exists j \in A^c \text{ s.t. } |X_j^\top (y - \bar{y} - \rho X_A \bar{w}_A)/s_j(\lambda)| = \lambda - \rho \right.$$

Stage 4
$$\begin{cases} \bar{\beta}_A(\lambda - \rho) = \bar{\beta}_A(\lambda) + \rho \bar{w}_A, \quad \bar{\beta}_{A^c}(\lambda - \rho) = 0, \quad \bar{y} = X\bar{\beta}(\lambda - \rho) \\ \lambda = \lambda - \rho \end{cases}$$

return $\bar{\beta}$

In the following we will compare of RepLars with the Lars by comparing RepLars with $\theta > 0$ and RepLars with $\theta = \mathbf{0}$. We will denote by $\hat{\beta}$ the Lars estimate corresponding to $\bar{\beta}$ produced by the algorithm above. Similarly, we let $\hat{w}_A$ be the vector in the Lars specialization corresponding to $\bar{w}_A$ above. We will assume throughout this analysis that variables are added to the active set one by one.

Theorem 4. *Assume that* **A1–2** *hold. Denote by* $\hat{\beta}(\lambda)$ *and* $\bar{\beta}(\lambda)$ *the Lars and RepLars solutions indexed by parameter* λ*. Conditioned on* X, y*, we have for any* $\lambda_{\min} > 0$

$$\forall \lambda \geq \lambda_{\min}\ S(\hat{\beta}(\lambda)) \subseteq S \implies \forall \lambda \geq \lambda_{\min}\ S(\bar{\beta}(\lambda)) \subseteq S. \tag{30}$$

Proof. Suppose that $\forall \lambda \geq \lambda_{\min}\ S(\hat{\beta}(\lambda)) \subseteq S$. Suppose further that $\hat{\beta}(\lambda_{\min})$ corresponds to iteration t of Lars. For $t' \leq t$ let $\hat{A}_{t'}$ be the active sets of Lars at iteration t'. With a slight abuse of notation we will temporarily treat an active set as an unordered set. Assumption **A2** guarantees that any variable that is at some point in the active set is also at some point in the support set. To see this, note that that by **A2**, the vector $\hat{w}_A$ never contains a zero element. If it did, then an equiangular vector u_A of X_A as in Eq. (2.6) of [3] could be constructed using a strict subset of vectors indexed by A, violating assumption **A2**. But if $\hat{w}_A$ does not contain a zero element, then the elements in the active set A cannot indefinitely be assigned a $\hat{\beta}_A(\lambda)$ coefficient of zero as λ is swept out. Finally, because we know $\forall \lambda \geq \lambda_{\min}, S(\hat{\beta}(\lambda)) \subseteq S$, this means that $\forall t' \leq t, \hat{A}_{t'} \subseteq S$. We will now argue by induction that the induced sequence of active sets $\bar{A}_{t'}$ of the RepLars also satisfies $\forall t' \leq t, \bar{A}_{t'} \subseteq S$.

Base case: Since $s(\|X^\top y\|_\infty) = \mathbf{1}$, the first variable selected by RepLars and the Lars method is the same. That is, $\hat{A}_1 = \bar{A}_1 \subseteq S$ at iteration 1.

Inductive step: Assume that $\forall t'' \leq t', \hat{A}_{t''} = \bar{A}_{t''} \subseteq S$. Since $\forall t'' \leq t', \bar{A}_{t''} \subseteq S$, we know by **A1** that for all λ up to iteration t', $s(\lambda)$ did not change on S, i.e. $s_S(\lambda) = \mathbf{1}$. This in particular means that both the Lars and the RepLars will have arrived at the same value of λ and intermediate estimate $\hat{\beta}(\lambda) = \bar{\beta}(\lambda)$ of β^* at the end of stage 4 of iteration $t' - 1$ and the same vectors $\hat{w}_A = \bar{w}_A$ at the end of stage 2 of iteration t'. To see this, notice that since $\forall t'' \leq t', \bar{A}_{t''} \subseteq S$, and since for all λ up to iteration t' we had $s_S(\lambda) = \mathbf{1}$, the RepLars is up to stage 2 of iteration t' equivalent to running Lars on the subset of variables X_S, y.

At stage 3 of iteration t', the RepLars algorithm determines which variable to add to $\bar{A}_{t'}$ in stage 1 of iteration $t' + 1$. Since the value of λ and the intermediate variables $\hat{\beta}(\lambda) = \bar{\beta}(\lambda)$ and $\hat{w}_A = \bar{w}_A$ are the same at stage 2 of iteration t' we can now use properties of $s(\lambda)$ to show that this implies $\hat{A}_{t'+1} = \bar{A}_{t'+1}$. Specifically, since (1) $s_S(\lambda) = \mathbf{1}$ did not change on S; (2) we always have $s(\lambda) \geq \mathbf{1}$; and (3) $\hat{A}_{t'+1} \subseteq S$, it follows that the RepLars will add the same variable. Hence, it follows that $\hat{A}_{t'+1} = \bar{A}_{t'+1} \subseteq S$.

By the principle of induction, we have shown that $\forall t' \leq t, \bar{A}_{t'} \subseteq S$. As the value of λ at the end of stage 4 of iteration t must be the same for RepLars and Lars, and since for that value we have by definition $\lambda < \lambda_{\min}$, we now know that $\forall \lambda \geq \lambda_{\min}, S(\bar{\beta}(\lambda)) \subseteq S$. □